Reasoners' Cooperation and Quantifier Elimination

Silvio Ghilardi
Technical Report n. 288-03
Dipartimento di Scienze dell'Informazione
Università degli Studi di Milano
Italy

March 20, 2003

Abstract

We exploit quantifier elimination in the global design of decision and semi-decision combined procedures for theories over non-disjoint signatures, thus providing in particular extensions of Nelson-Oppen results.

Keywords. Combination, Nelson-Oppen Procedure, Fusion, Superposition Calculus, Quantifier Elimination, Model Completions.

1 Introduction

Quantifier elimination has been considered, since the early times of modern symbolic logic, a powerful technique for decision procedures. Even in actual approaches to combination problems (see e.g. [12]), specific quantifier elimination algorithms are often invoked as specialized reasoners to be integrated within a flexible general setting dealing with multiple theories. This happens, in particular, whenever numerical constraints problems [5] need to be adequately addressed: examples of such specialized reasoners are for instance the Fourier-Motzkin quantifier elimination procedure (and its refinements [23]) for linear rational arithmetic or Cooper's [7], [16] quantifier elimination procedure for integer Presburger arithmetic.

In contrast to this *local call* for quantifier elimination algorithms, we shall address in this paper quantifier elimination as a *global design* opportunity for integrated provers in the Nelson-Oppen style.

To see what we mean in a rather simple example, let us consider Nelson-Oppen original combination procedure [14], [17], [20]: here we are given theories T_1, T_2 over disjoint signatures (with equality) Σ_1, Σ_2 and we are dealing with the problem of deciding the universal fragment of the union theory $T_1 \cup T_2$ (given that the universal fragments of T_1 and T_2 are decidable).[1] Taking into consideration conjunctive normal forms, after the necessary skolemization and purification steps, the problem reduces to that of establishing the consistency of sets of sentences like

$$(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2) \tag{1}$$

where Γ_i $(i = 1, 2)$ are finite sets of ground literals in the signature Σ_i augmented with new *shared* Skolem constants, say $\underline{a} = a_1, \dots, a_n$. Let us call Σ_0 the common signature $\Sigma_1 \cap \Sigma_2$ (which is reduced to the only equality predicate). It is evident that, if (1) has a model $\mathcal{M}$, then we can take the Robinson diagram $\Delta(\mathcal{A})$ of the Σ_0-substructure $\mathcal{A}$ of $\mathcal{M}$ generated by $\underline{a}$ and realize that the two set of sentences

$$T_1 \cup \Gamma_1 \cup \Delta(\mathcal{A}) \qquad \text{and} \qquad T_2 \cup \Gamma_2 \cup \Delta(\mathcal{A}) \tag{2}$$

are consistent as well, because $\mathcal{M}$ is a model of both of them. $\Delta(\mathcal{A})$ simply specifies exhaustively whether for any $i, j = 1, \dots, n$, we have $\mathcal{M} \models a_i = a_j$ or $\mathcal{M} \models a_i \neq a_j$. Since there are only finitely many such 'arrangements' $\Delta(\mathcal{A})$, it is clear that if the consistency of the two sets (2) is not only a necessary but also a sufficient condition for the consistency of (1), then we actually have a combined decision procedure.

The consistency of both the sets in (2) is not precisely sufficient for the consistency of (1), however it is sufficient whenever T_1, T_2 are *stably infinite*: this means that every model of T_i $(i = 1, 2)$ embeds into an infinite model of T_i.[2] Otherwise said, we have the following fact (here quantifier elimination comes into the picture): the shared universal theory T_0 (which is nothing but the pure equality theory) can be extended to a theory T_0^* in the same signature Σ_0 (namely the theory of an infinite set) which is conservative over the universal fragment of T_0 and which, in addition, enjoys quantifier elimination. In the standard (robinsonian) model-theoretic terminology, T_0^* is a *model completion* of T_0 and stable infiniteness means that every model of T_i $(i = 1, 2)$ embeds into a model of $T_i \cup T_0^*$.

This extra hypothesis is indeed sufficient for the consistency of the two sets of sentences in (2) to imply the consistency of the set of sentences in (1), by a simple proof based on

[1] Here by 'deciding the universal fragment of a theory T' we mean to decide whether a quantifier free formula ϕ is a logical consequence of T. As a formula is a logical consequence of a theory T iff its negation is T-unsatisfiable, a decision procedure for the universal fragment of T is equivalent to a decision procedure for T-satisfiability of constraints which are expressible as existential closures of quantifier free formulas.

[2] The usual definition of stable infiniteness required by the Nelson-Oppen procedure is not just this one, but it is equivalent to it by a simple compactness argument, see Proposition 3.2 below.

submodel completeness and on Robinson's joint consistency theorem (see section 5 for the details).

What does this argument need to work? It needs that: i) there is a universal theory T_0 in the shared signature Σ_0 which is contained in both T_1 and T_2; ii) T_0 admits a model-completion T_0^*; iii) every model of T_i embeds into a model of $T_i \cup T_0^*$; iv) T_0 is locally finite (this last point is needed for having finitely many arrangements to try with).

This analysis makes clear that our argument is quite general and extends Nelson-Oppen procedure to non-disjoint cases: we shall use it in section 6 to obtain results concerning fusion decidability [24] in propositional modal logic.

Whenever only hypothesis iv) fails, we no longer have a combined decision procedure; however, even in this case, the method can be used in order to limit partial residue exchange to disjunctions of positive literals (see section 4 below). Moreover, it can be applied in saturation-based theorem-proving in order to forbid inferences with mixed languages premises, still retaining refutational completeness (see section 7 below).

We shall begin the paper with a section recalling the basic model-theoretic ingredients in more detail and providing also examples from natural software verification theories.

2 Model-theoretic background

A *signature* Σ is a set of functions and predicate symbols (each of them endowed with the corresponding arity). We assume the binary equality predicate symbol $=$ to be always present in Σ.[3] The signature obtained from Σ by the addition of a set of new constants ($=$ 0-ary function symbols) X is denoted by $\Sigma \cup X$ or by Σ^X. We have the usual notions of Σ-*term*, (full first order) -*formula*, -*atom*, -*literal*, -*clause*, -*positive clause*, etc.: e.g. atoms are just atomic formulas, literals are atoms and their negations, clauses are multisets of literals, positive clauses are multisets of atoms, etc. Whenever confusion does not arise, we may improperly write clauses/positive clauses as disjunctions of the corresponding literals/atoms (in particular, the empty clause is always confused with the empty disjunction $\perp$ expressing syntactic falsity). Letters $\phi, \psi, \ldots$ are used for formulas, whereas letters $A, B, \ldots$ are used for literals and letters $C, D, \ldots$ are used for clauses. Terms, literals, clauses and formulas are called *ground* whenever free variables do not appear in them. Ground formulas are also called sentences. A Σ-*theory* T is a set of sentences (called the axioms of T) in the signature Σ; however when we write $T \subseteq T'$ for theories, we may mean not just set-theoretic inclusion but the fact that all the axioms

[3]In order to avoid confusion, we use in the metalanguage the symbol $\equiv$ (instead of $=$) to mean identity of syntactic expressions.

for T are logical consequences of the axioms for T'. We occasionally use also the standard automatic deduction notations $E_{|p}, E[s]_p$ to denote the subterm at position p in the expression (=term, atom, literal, clause, etc.) E and the expression resulting from E by the replacement of subterm at position p by the term s.

From the semantic side, we have the standard notion of a Σ-*structure* $\mathcal{A}$: this is nothing but a support set endowed with an arity-matching interpretation of the predicate and function symbols from Σ (we use $f^{\mathcal{A}}$ to denote the interpretation of the function symbol f in the structure $\mathcal{A}$). We shall notationally confuse, for the sake of simplicity, a structure with its support set. Truth of a Σ-formula in $\mathcal{A}$ is defined in any one of the standard ways (so that truth of a formula is equivalent to truth of its *universal* closure). A Σ-structure $\mathcal{A}$ is a *model* of a Σ-theory T (in symbols $\mathcal{A} \models T$) iff all axioms of T are true in $\mathcal{A}$; for models of a Σ-theory T we shall preferably use the letters $\mathcal{M}, \mathcal{N}, \dots$ to distinguish them from arbitrary Σ-structures. If ϕ is a formula, $T \models \phi$ ('ϕ *is a logical consequence of* T') means that ϕ is true in any model of T. A Σ-theory T is *complete* iff for every Σ-sentence ϕ, either ϕ or $\neg\phi$ is a logical consequence of T; T is *consistent* iff it has a model (i.e. iff $T \not\models \bot$).

The main problems we deal with are *word problems*, more precisely, given a Σ-theory T:

- the *word problem* for T is that of deciding whether $T \models A$ holds for a Σ-atom A;
- the *conditional word problem* for T is that of deciding whether $T \models C$ holds for a Σ-clause C containing at most one atom;
- the *clausal word problem* for T is that of deciding whether $T \models C$ holds for a Σ-clause C;
- the *elementary word problem* for T is that of deciding whether $T \models \phi$ holds for a first order Σ-formula ϕ.

An *embedding* between two Σ-structures $\mathcal{A}$ and $\mathcal{B}$ is any map $f : \mathcal{A} \longrightarrow \mathcal{B}$ among the corresponding support sets satisfying the condition

$$(*) \qquad \mathcal{A} \models A \quad \text{iff} \quad \mathcal{B} \models A$$

for all $\Sigma^{\mathcal{A}}$ atoms A (here $\mathcal{A}$ is regarded as a $\Sigma^{\mathcal{A}}$-structure by interpreting each $a \in \mathcal{A}$ into itself and $\mathcal{B}$ is regarded as a $\Sigma^{\mathcal{A}}$-structure by interpreting each $a \in \mathcal{A}$ into $f(a)$). Notice that, as we have identity in the language, an embedding is an injective function (it also must preserve the interpretation of the function symbols and, in case it is just an inclusion, the interpretation of the predicate symbols in the smaller structure must be just the restriction of the corresponding interpretation in the bigger structure). In case $(*)$ holds for all first order formulas, the embedding is said to be *elementary*.

The *diagram* $\Delta(\mathcal{A})$ of a Σ-structure $\mathcal{A}$ is the set of ground $\Sigma^{\mathcal{A}}$-literals which are true in $\mathcal{A}$; the elementary diagram $\Delta^e(\mathcal{A})$ of a Σ-structure $\mathcal{A}$ is the set of $\Sigma^{\mathcal{A}}$-sentences which are true in $\mathcal{A}$. Robinson (elementary) diagram theorem [6] says that there is an (elementary) embedding between the Σ-structures $\mathcal{A}$ and $\mathcal{B}$ iff it is possible to expand $\mathcal{B}$ to a $\Sigma^{\mathcal{A}}$-structure in such a way that it becomes a model of the (elementary) diagram of $\mathcal{A}$. This theorem (in combination with the compactness of the logical consequence relation) will be repeatedly used without explicit mention in the paper; a typical standard use is the following: suppose that we want to embed $\mathcal{A}$ into a model of a theory T, then it is sufficient to check that $T \cup \Delta(\mathcal{A})$ is consistent.

A formula is *open* or quantifier-free iff it does not contain quantifiers. A Σ-theory T is said to *eliminate quantifiers* iff for every formula $\phi(\underline{x})$[4] there is an open formula $\phi'(\underline{x})$ such that

$$T \models \phi(\underline{x}) \leftrightarrow \phi'(\underline{x}).$$

There are many well-known theories [6] eliminating quantifiers, we give here some examples which can be of interest for software verification.

Example. Linear integer arithmetic (i.e. the theory of the structure of integer numbers in the signature $+, 0, 1, \leq, \equiv_n$) eliminates quantifiers; so does rational linear arithmetic (i.e. the theory of rational numbers in the signature $+, 0, \leq$). Another well-known classical example from Tarski is real arithmetic (i.e. the theory of real numbers in the signature $+, 0, \cdot, 1, \leq$). $\dashv$

Example. The theory of acyclic binary lists L [17], [18] eliminates quantifiers (see Appendix A): this is the theory in the signature

$$\Sigma_L = \{car, cdr, cons\}$$

consisting on the universal closures of the following axioms:

$$\begin{aligned}
&cons(car(x), cdr(x)) = x\\
&car(cons(x,y)) = x\\
&cdr(cons(x,y)) = y\\
&x \neq t(x)
\end{aligned}$$

[4]By this notation, we mean that ϕ contains free variables only among the finite set $\underline{x}$.

where t is a term built up from x by using finitely many applications of the unary function symbols car, cdr. Notice that all models of L are infinite: in fact, by the first three axioms, the interpretation of the symbol $cons$ realizes, in any model $\mathcal{M}$ of L, a bijection $\mathcal{M} \times \mathcal{M} \longrightarrow \mathcal{M}$ and such a bijection cannot exist if the support of $\mathcal{M}$ is a finite set having more than one element (it cannot have just one element because $L \models \forall x\,(car(x) \neq x)$).[5] $\dashv$

The main ingredient of this paper is the well-known notion of a *model completion* of a theory. There are good chapters on that in all texbooks from Model Theory. We shall recall here just the essential definitions *for the only case of universal theories*[6] which is the relevant one for the purposes of this paper (but see Appendix B). Readers may consult e.g. [13],[6], [22] for further information.

Let T be a universal Σ-theory and let $T^* \supseteq T$ a further Σ-theory; we say that T^* is a model completion of T iff i) every model of T has an embedding into a model of T^* and ii) T^* eliminates quantifiers.

It can be shown that a model completion T^* of a theory T is unique, in case it exists, and moreover that T^* has a set of $\forall\exists$-axioms,[7] see [6].

Example. The theory of an infinite set is the model completion of pure equality theory; the theory of dense total orders without endpoints is the model completion of the theory of total orders. $\dashv$

Example. There are many classical examples from algebra: the theory of algebraically closed fields is the model completion of the theory of intergral domains, the theory of divisible torsion free abelian groups is the model completion of the theory of torsion free abelian groups, etc. $\dashv$

Example. The theory of atomless Boolean algebras is the model completion of the theory of Boolean algebras (for model completions arising in the algebra of logic, see the book [11]). We recall that an atom in a Boolean algebra $\mathcal{B}$ is a non-zero element which is minimal (among non-zero elements) with respect to the partial order induced by the lattice structure. A Boolean algebra $\mathcal{B}$ is atomless iff it has no atoms. $\dashv$

[5] It is customary to assume that the support of a structure cannot be the empty set. This restriction is unjustified and, according to common folklore for people belonging to the categorical logic community, it can be removed simply by marking the free variables in a syntactic derivation. In case empty supports are allowed, however, in order to get quantifier elimination for L, one should add either a free constant to the language or the further axiom $\exists x(x = x)$ (over the empty string of variables!) to the theory.

[6] Recall that a universal theory T is a theory having as axioms only universal closures of open formulas.

[7] A sentence is $\forall\exists$ iff it is obtained from an open formula by prefixing it some existential quantifiers and finally some universal quantifiers (in this order).

Example. An old result in [22] says, in particular, that universal Horn theories T in finite signatures always have a model completion, provided the following two conditions are satisfied: a) finitely generated models of T are all finite; b) amalgamation property holds for models of T. This fact can be used in order to prove the existence of a model completion for theories axiomatizing many interesting discrete structures (like graphs, posets, etc.). $\dashv$

Example. It follows from the quantifier elimination results of Appendix A that the theory L of acyclic lists is the model completion of itself. $\dashv$

Example. If a theory T^* has elimination of quantifiers, then it is the model completion of the theory T axiomatized by the set of universal sentences which are logical consequences of T, see [6]. $\dashv$

3 Compatibility

The key ingredient for our combination procedures is the following notion:

Definition 3.1 *Let T be a theory in the signature Σ and let T_0 be a universal theory in a subsignature $\Sigma_0 \subseteq \Sigma$. We say that T is* T_0-compatible *iff*

(i) $T_0 \subseteq T$;

(ii) *T_0 has a model-completion T_0^*;*

(iii) *every model of T embeds into a model of $T \cup T_0^*$.*

Condition (iii) can be equivalently given in a slightly different form:

Proposition 3.2 *Let T be a theory in the signature Σ and let $T_0 \subseteq T$ be a universal theory in a subsignature $\Sigma_0 \subseteq \Sigma$. If T_0 has a model-completion T_0^*, then T is T_0-compatible iff every quantifier-free Σ-formula which is false in a model of T is false also in a model of $T \cup T_0^*$.*[8]

Proof. It is evident that if every model of T embeds into a model of $T \cup T_0^*$, then every quantifier-free formula in Σ which fails in a model of T fails in a model of $T \cup T_0^*$

[8]Equivalently, iff every existential closure of an open Σ-formula which is satisfiable in a model of T is satisfiable also in a model of $T \cup T_0^*$. To understand the statement of Proposition, recall that truth of a formula in a structure is equivalent to truth of its universal closure.

too. The converse is by a compactness argument: let $\mathcal{M}$ be a model of T, we show that $T \cup T_0^* \cup \Delta(\mathcal{M})$ is consistent. If not, we have $T \cup T_0^* \models \neg\phi(\underline{a})$ where $\phi(\underline{a})$ is a quantifier-free ground formula (with parameters $\underline{a}$ from the support of $\mathcal{M}$), which is true in $\mathcal{M}$. Abstracting variables from parameters, we get $T \cup T_0^* \models \forall \underline{x} \neg\phi(\underline{x})$. But $\neg\phi(\underline{x})$ is a quantifier-free formula failing in a model of T and, by hypothesis, there is a model of $T \cup T_0^*$ in which $\neg\phi(\underline{x})$ fails too, contrary to the fact that $T \cup T_0^* \models \forall \underline{x} \neg\phi(\underline{x})$. $\dashv$

Example According to the above Proposition, it is evident that T_0-compatibility reduces to the standard notion of stable infiniteness (used in the disjoint Nelson-Oppen combination procedure) in case T_0 is the pure theory of equality. $\dashv$

Example Every theory including the theory L of acyclic lists is compatible with L, because L is universal and $L = L^*$. $\dashv$

Example If T_0 has a model completion T_0^* and if $T \supseteq T_0^*$, then T is certainly T_0-compatible: this trivial case is often interesting (we may take e.g. T_0 to be the theory of linear orders and T to be real aritmetic or rational linear arithmetic). $\dashv$

Example. Let T_0 be a universal theory having a model completion T_0^* and let T be any extension of T_0 with free function symbols only. In this case, T is T_0-compatible: to see it, take any model $\mathcal{M}$ of T, embeds its Σ_0-reduct into a model $\mathcal{M}'$ of T_0^* and expand in any arbitrary way the interpretation of the free function symbols to the tuples of $\mathcal{M}'$ not entirely belonging to $\mathcal{M}$. $\dashv$

More examples will be supplied in section 6. An interesting feature of T_0-compatibility is that it is a *modular* property; we formally state this fact now and prove it in section 5:

Proposition 3.3 *Let T_1 be a Σ_1-theory and let T_2 be a Σ_2-theory ; suppose they are both compatible with respect to a Σ_0-theory T_0 (where $\Sigma_0 := \Sigma_1 \cap \Sigma_2$). Then $T \cup T_2$ is T_0-compatible too.*

4 Combining compatible theories

Let us progressively fix our main data for the whole paper.

Assumption (I). *T_1 is a theory in the signature Σ_1 and T_2 is a theory in the signature Σ_2; Σ_0 is the signature $\Sigma_1 \cap \Sigma_2$.*

Our main aim is that of (semi)deciding the clausal word problem for $T_1 \cup T_2$, given that the corresponding clausal word problems for T_1 and T_2 are (semi)decidable. By skolemization, this amounts to (semi)decide the consistency of

$$T_1 \cup T_2 \cup \Gamma, \tag{1}$$

where Γ is a finite set of ground literals in the signature $\Sigma_1 \cup \Sigma_2$, expanded with a finite set of new Skolem constants.

Γ can be *purified.* In fact, for any literal $A \in \Gamma$, we have that (1) is equi-satisfiable with

$$T_1 \cup T_2 \cup (\Gamma \setminus \{A\}) \cup \{A[c]_p,\ c = A_{|p}\}, \tag{2}$$

where p is a term position in A and c is a further new constant. After finitely many transformations like that from (1) to (2), we can reduce our problem to that of establishing the consistency of a set of sentences like

$$(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2), \tag{3}$$

where Γ_1, Γ_2 are as explained in the following:[9]

Assumption (II). *For finitely many new free constants $\underline{a}$, Γ_1 is a finite set of ground literals in the signature $\Sigma_1^{\underline{a}}$ and Γ_2 is a finite set of ground literals in the signature $\Sigma_2^{\underline{a}}$.*

Clearly the consistency of (3) cannot follow from the mere separate consistency of $T_1 \cup \Gamma_1$ and of $T_2 \cup \Gamma_2$ (for trivial reasons, take e.g. $T_1 = T_2 = \emptyset$, $\Gamma_1 = \{a_1 = a_2\}$ and $\Gamma_2 = \{a_2 = a_3,\ a_1 \neq a_3\}$). We need some *information exchange* between a reasoner dealing with $T_1 \cup \Gamma_1$ and a reasoner dealing with $T_2 \cup \Gamma_2$.

Craig's interpolation theorem for first order logic ensures that the inconsistency of (3) can be detected by the information exchange of a single $\Sigma_0^{\underline{a}}$-sentence ϕ such that $T_1 \cup \Gamma_1 \models \phi$ and $T_2 \cup \Gamma_2 \cup \{\phi\} \models \bot$. However, as pointed out in [19], this observation is not very useful, as ϕ might be any first-order formula, whereas we would like - at least - ϕ to be quantifier-free: remind that most existing provers detect inconsistency just by

[9] We leave the reader the details of this standard step. Of course, one can apply the above transformation until in each literal there is at most one symbol which is not the identity predicate or a Skolem constant, but this method (although correct) would be highly inefficient. A better solution, which is commonly adopted, is that of abstracting only 'alien' subterms; however, there are many different possibilities for defining what an alien subterm is in case the signatures are not disjoint (we leave the reader to consult [3] for a detailed discussion). Notice that at the end of this purification process, Σ_0-literals could be inserted either in Γ_1 or in Γ_2 or in both of them, indifferently.

skolemization and saturation, so they surely wouldn't be able to find such a ϕ (if it is not quantifier-free), even in case they can efficiently handle with both $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$.

Unfortunately, information exchange of $\Sigma_0^{\underline{a}}$-quantifier free formulas alone is not sufficient, even for syntactically simple T_1 and T_2, to establish the inconsistency of (3). An example, with $T_1 =$ the theory of Boolean algebras and $T_2 =$ the theory of pseudocomplemented distributive lattices, will be provided in section 7 below. We so need a further assumption in order to get limited information echange without affecting refutational completeness (this is the relevant assumption we make, the other two being mere notational conventions):

Assumption (III). *There is a universal Σ_0-theory T_0 such that both T_1 and T_2 are T_0-compatible.*

A finite list

$$C_1, \dots, C_n$$

of positive ground $\Sigma_0^{\underline{a}}$-clauses such that for every $k = 1, \dots, n$, there is $i = 1, 2$ such that

$$T_i \cup \Gamma_i \cup \{C_1, \dots, C_{k-1}\} \models C_k.$$

is called a *positive residue chain.* We can now formulate our main combination results, whose proof can be found in the next section:

Theorem 4.1 *In the above assumptions, $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent iff there is a positive residue chain $C_1, \dots, C_n$ such that C_n is the empty clause.*

Thus inconsistency can be detected by repeated exchanges of positive ground clauses only; if we allow information exchange consisting on ground quantifier free formulas, a single exchange step is sufficient:

Theorem 4.2 *In the above assumptions, $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent iff there is a ground quantifier-free $\Sigma_0^{\underline{a}}$-sentence ϕ such that*

$$T_1 \cup \Gamma_1 \models \phi \qquad \textit{and} \qquad T_2 \cup \Gamma_2 \cup \{\phi\} \models \bot.$$

We recall that our T_i's are said to be Σ_0-convex [19] iff whenever it happens that $T_i \cup \Gamma_i \models A_1 \vee \cdots \vee A_n$ (for $n \geq 1$ and for ground $\Sigma_0^{\underline{a}}$-atoms $A_1, \dots, A_n$), then there is $k = 1, \dots, n$ such that $T_i \cup \Gamma_i \models A_k$.[10] For Σ_0-convex theories, an immediate subsumption argument refines Theorem 4.1 in the following way:

[10] Among Σ_0-convex theories we have the important class of universal Horn theories, see [19] again.

Corollary 4.3 *In addition to the above assumptions, suppose also that T_1, T_2 are both Σ_0-convex. Then $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent iff there is a positive residue chain $C_1, \ldots, C_n$ in which $C_1, \ldots, C_{n-1}$ are all ground Σ_0-atoms and C_n is $\perp$.*

5 Proofs of the combination results

This section is devoted only to the proofs of Theorems 4.1, 4.2 and of Proposition 3.3.

Say that a set Γ_0 of *positive* ground $\Sigma_0^{\underline{a}}$-clauses is *saturated* iff it is closed under the two rules

$$T_1 \cup \Gamma_1 \cup \Gamma_0 \models C \qquad \Rightarrow \qquad C \in \Gamma_0$$

$$T_2 \cup \Gamma_2 \cup \Gamma_0 \models C \qquad \Rightarrow \qquad C \in \Gamma_0.$$

Lemma 5.1 *Suppose that Γ_0 is saturated and does not contain the empty clause. Then there are $\Sigma_i^{\underline{a}}$-models $\mathcal{M}_i$ $(i = 1, 2)$ such that*

$$\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0 \quad \textit{and} \quad \mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0;$$

moreover $\mathcal{M}_1$ and $\mathcal{M}_2$ share the same Σ_0-substructure generated by the elements (denoted by) $\underline{a}$.

Proof. The statement of the lemma is proved if we are able to find an exhaustive[11] set Δ of ground $\Sigma_0^{\underline{a}}$-literals which is consistent with both $T_1 \cup \Gamma_1 \cup \Gamma_0$ and $T_2 \cup \Gamma_2 \cup \Gamma_0$. In this case, in fact, given any two models $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0 \cup \Delta$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0 \cup \Delta$, we have that their Σ_0-substructures generated by $\underline{a}$ both have diagram Δ, consequently they are Σ_0-isomorphic (hence we may assume that they are just the same, by renaming some elements in one of the supports, if needed).

We shall adapt the notion of productive clause used in nowadays refutational completeness proofs for e.g. resolution or paramodulation based calculi. Consider any strict total terminating order on ground $\Sigma_0^{\underline{a}}$-atoms and extend it to a strict total terminating order $>$ for positive ground $\Sigma_0^{\underline{a}}$-clauses by taking standard multiset extension. We shall define increasing sets Δ_C^+ (varying $C \in \Gamma_0$) of ground $\Sigma_0^{\underline{a}}$-atoms as follows. Recall that, as the empty clause is not in Γ_0, all positive clauses in Γ_0 are of the kind $A \vee A_1 \vee \cdots \vee A_n$ $(n \geq 0)$.

The definition is by transfinite induction on $>$. Say that the clause $C \equiv A \vee A_1 \vee \cdots \vee A_n$ from Γ_0 is *productive* iff i) $\{A\} > \{A_1, \ldots, A_n\}$ and ii) $A_1, \ldots, A_n \notin \Delta_{<C}^+$ (where $\Delta_{<C}^+$ is

[11] This means that Δ must contain, for every ground $\Sigma_0^{\underline{a}}$-literal A, either A itself or its negation.

$\bigcup_{D<C} \Delta^+_D$). Now, if C is productive, we let Δ^+_C to be $\Delta^+_{<C} \cup \{A\}$, otherwise Δ^+_C is simply $\Delta^+_{<C}$.

Let Δ^+ be $\bigcup_{C\in\Gamma_0} \Delta^+_C$ and Δ be $\Delta^+ \cup \{\neg A \,|\, A$ is a ground $\Sigma^{\underline{a}}_0$-atom not belonging to $\Delta^+\}$. By construction, $\Delta \models \Gamma_0$, so we simply need to show that $T_1 \cup \Gamma_1 \cup \Delta$ and $T_2 \cup \Gamma_2 \cup \Delta$ are consistent. We need a preliminary claim.

Claim: if the clause $A \vee A_1 \vee \cdots \vee A_n$ is productive and A is the maximum atom in it, then $A_1, \ldots, A_n \not\in \Delta^+$: this is evident, as the A_i's could only be produced by clauses smaller than $A \vee A_1 \vee \cdots \vee A_n$.

Suppose now that $T_1 \cup \Gamma_1 \cup \Delta$ is not consistent (the case $i = 2$ is analogous). Then there are ground atoms $B_1, \ldots, B_m \not\in \Delta^+$ and productive clauses

$$\begin{aligned} C_1 &\equiv \quad A_1 \vee A_{11} \vee \cdots \vee A_{1k_1} \\ &\quad\ \ \ldots \\ C_n &\equiv \quad A_n \vee A_{n1} \vee \cdots \vee A_{nk_n} \end{aligned}$$

(with maximum atoms $A_1, \ldots, A_n$, respectively), such that

$$T_1 \cup \Gamma_1 \cup \{A_1, \ldots, A_n\} \models B_1 \vee \cdots \vee B_m.$$

By trivial logical manipulations, it follows that

$$T_1 \cup \Gamma_1 \cup \{C_1, \ldots, C_n\} \models \bigvee_{i,j} A_{ij} \vee B_1 \vee \cdots \vee B_m.$$

As $C_1, \ldots, C_n$ are clauses in Γ_0 and as Γ_0 is saturated, the clause

$$D \equiv \bigvee_{i,j} A_{ij} \vee B_1 \vee \cdots \vee B_m$$

is also in Γ_0. By construction (anyway, either D is productive or not) some of the atoms $\{A_{11}, \ldots, A_{nk_n}, B_1, \ldots, B_m\}$ is in Δ^+. By the claim, $A_{11}, \ldots, A_{nk_n}$ cannot be there, so one of the B_j's is in Δ^+, contradiction. $\dashv$

If $\mathcal{M}$ is a Σ-model and $X \subseteq \mathcal{M}$, we can consider $\mathcal{M}$ as a $\Sigma\cup X$-structure by interpreting the name $\bar{b}$ of each $b \in X$ into b. Next lemma uses the assumption on T_0-compatibility:

Lemma 5.2 *Let $\mathcal{M}_1$ be a Σ_1-model of T_1 and let $\mathcal{M}_2$ be a Σ_2-model of T_2; suppose also that $\mathcal{M}_1$ and $\mathcal{M}_2$ share a common Σ_0-substructure $\mathcal{A}$. Then there are a $\Sigma_1\cup\Sigma_2\cup\mathcal{A}$-model $\mathcal{M}$ of $T_1 \cup T_2$ and two $\Sigma_i^{\mathcal{A}}$-embeddings $\mathcal{M}_i \longrightarrow \mathcal{M}$ $(i = 1, 2)$.*

Proof. By T_0-compatibility, we can suppose that the $\mathcal{M}_i$ are models of $T_i \cup T_0^*$. By renaming some elements in the supports if needed, we can also freely suppose that the sets

$\mathcal{M}_1 \setminus \mathcal{A}$ and $\mathcal{M}_2 \setminus \mathcal{A}$ are disjoint. Consider the elementary diagrams $\Delta^e(\mathcal{M}_1)$, $\Delta^e(\mathcal{M}_2)$ of $\mathcal{M}_1$, $\mathcal{M}_2$; we show that $\Delta^e(\mathcal{M}_1) \cup \Delta^e(\mathcal{M}_2)$ is consistent as a $\Sigma_1 \cup \Sigma_2 \cup \mathcal{M}_1 \cup \mathcal{M}_2$-theory.

First notice that $T_0^* \cup \Delta(\mathcal{A})$ is a complete $\Sigma_0^{\mathcal{A}}$-theory: this is just submodel-completeness, namely the semantic characterization of quantifier elimination. We are actually using the trivial direction of such a characterization: the reason why $T_0^* \cup \Delta(\mathcal{A})$ is complete is the following. Consider two models $\mathcal{N}_1, \mathcal{N}_2$ of its and a $\Sigma_0^{\mathcal{A}}$-sentence $\phi(\underline{a})$. Now (by Robinson diagram theorem) $\mathcal{A}$ is a common substructure of $\mathcal{N}_1$ and $\mathcal{N}_2$; moreover $\phi(\underline{x})$ is T_0^*-equivalent to a quantifier free formula $\phi'(\underline{x})$, hence if $\phi(\underline{a})$ is true in $\mathcal{N}_1$, $\phi'(\underline{a})$ is true in it, $\phi'(\underline{a})$ is true in $\mathcal{A}$ and in $\mathcal{N}_2$, as well,[12] thus establishing that $\phi(\underline{a})$ is true in $\mathcal{N}_2$. Since $\phi(\underline{a})$, $\mathcal{N}_1, \mathcal{N}_2$ are arbitrary, $T_0^* \cup \Delta(\mathcal{A})$ is complete.

Having established that $T_0^* \cup \Delta(\mathcal{A})$ is a complete theory, we see that $\Delta^e(\mathcal{M}_1)$ and $\Delta^e(\mathcal{M}_2)$ are both extensions of its. Now we can simply invoke Robinson's joint consistency theorem: this theorem (see [6]) says that if H_0 is a complete theory in a signature $\Theta_0 = \Theta_1 \cap \Theta_2$ and H_1, H_2 are consistent extensions of its in the signatures Θ_1, Θ_2, respectively, then $H_1 \cup H_2$ is consistent too in the signature $\Theta_1 \cup \Theta_2$ (in our case, Θ_1 is $\Sigma_1 \cup \mathcal{M}_1$ and Θ_2 is $\Sigma_2 \cup \mathcal{M}_2$, so that Θ_0 is $\Sigma_0 \cup \mathcal{A}$, because the sets $\mathcal{M}_1 \setminus \mathcal{A}$ and $\mathcal{M}_2 \setminus \mathcal{A}$ are disjoint). Having established that $\Delta^e(\mathcal{M}_1) \cup \Delta^e(\mathcal{M}_2)$ is consistent, clearly any model of its fits our $\mathcal{M}$ (notice that such an $\mathcal{M}$ is a model of T_1 and T_2, because it is a model of the elementary diagrams of $\mathcal{M}_1$ and $\mathcal{M}_2$). $\dashv$

Proof of Theorem 4.1. Suppose that there is no positive residue chain ending up with the empty clause. We build a saturated set Γ_0 of of positive ground $\Sigma_0^{\underline{a}}$-clauses in ω steps. Let Θ_0 be the empty set; if Θ_k has already been defined, let Θ_{k+1} be the set of positive ground $\Sigma_0^{\underline{a}}$-clauses C such that $T_i \cup \Gamma_i \cup \Theta_k \models C$ holds for $i = 1$ or $i = 2$. Clearly $\Theta_k \subseteq \Theta_{k+1}$; moreover, by the compactness theorem for first order logic, it is clear that $\Gamma_0 = \bigcup_k \Theta_k$ is saturated. Notice also that a clause C belongs to Θ_{k+1} ($k \geq 0$) iff there is a positive residue chain $C_1, \dots, C_n, C$ such that $C_1, \dots, C_n$ all belong to Θ_k (this is easily proved by induction on k and by compactness again).[13] Consequently, Γ_0 does not contain the empty clause and Lemma 5.1 applies. This means that there are models $\mathcal{M}_1 \models T_1 \cup \Gamma_1$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2$ whose Σ_0-substructures generated by $\underline{a}$ are the same. We can now apply Lemma 5.2 to $\mathcal{M}_1, \mathcal{M}_2, \mathcal{A}$ (where $\mathcal{A}$ is this Σ_0-substructure generated by $\underline{a}$). By that Lemma, there are a model $\mathcal{M} \models T_1 \cup T_2$ and $\Sigma_i^{\mathcal{A}}$-embeddings $\mathcal{M}_i \longrightarrow \mathcal{M}$. As $\mathcal{M}_i \models \Gamma_i$, we have also $\mathcal{M} \models \Gamma_i$ ($i = 1, 2$) (recall that the Γ_i's are sets of ground

[12] Recall that ground open formulas are preserved by both sub- and super-structures.

[13] The induction step is as follows: if $C \in \Theta_{k+1}$, then there are $C_1, \dots, C_n \in \Theta_k$ such that $T_i \cup \Gamma_i \cup \{C_1, \dots, C_n\} \models C$ holds for $i = 1$ or $i = 2$. Now it is sufficient to append C to any juxtaposition of positive residue chains ending up in $C_1, \dots, C_n$.

$\Sigma_0^{\underline{a}}$-literals, so their truth is preserved by $\Sigma_i^{\mathcal{A}}$-embeddings). Thus $\mathcal{M} \models T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$, so the latter set is indeed consistent. $\dashv$

Proof of Theorem 4.2. We reduce this Theorem to the previous one. If $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent, there is a positive residue chain $C_1, \dots, C_n$ ending up with the empty clause. Say that C_k is an i-residue ($i = 1, 2$) iff $T_i \cup \Gamma_i \cup \{C_1, \dots, C_{k-1}\} \models C_k$. Let ψ_k (for $k = 1, \dots, n$) be the quantifier-free ground $\Sigma_0^{\underline{a}}$-formula $\neg C_1 \vee \cdots \vee \neg C_{k-1} \vee C_k$ and let ϕ be the conjunction of the ψ_k such that C_k is a 1-residue. Clearly, $T_1 \cup \Gamma_1 \models \phi$. Moreover, by induction, it is easy to see that $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ for all $j = 1, \dots, n$: in fact, if C_j is a 2-residue, then $T_2 \cup \Gamma_2 \cup \{C_1, \dots, C_{j-1}\} \models C_j$ (hence $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ by induction hypothesis) and if C_j is a 1-residue, then $\{\phi\} \models \neg C_1 \vee \cdots \vee \neg C_{j-1} \vee C_j$ (hence $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ again by induction hypothesis). As $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ holds for all j, we have in particular that $T_2 \cup \Gamma_2 \cup \{\phi\} \models \perp$ for $j = n$. $\dashv$

The same argument used for the proof of Lemma 5.2 gives also the

Proof of Proposition 3.3. Take a model $\mathcal{M}$ of $T_1 \cup T_2$ and embeds its Σ_i-reducts into models $\mathcal{M}_i$ of $T_i \cup T_0^*$ ($i = 1, 2$). As in the proof of Lemma 5.2, we can freely suppose that the embeddings are inclusions and that we have $\mathcal{M} = \mathcal{M}_1 \cap \mathcal{M}_2$ for supports. Now $T_0^* \cup \Delta(\mathcal{M})$ is a complete theory (here $\Delta(\mathcal{M})$ is the diagram of $\mathcal{M}$ as a Σ_0-structure), hence by Robinson consistency theorem there is a model $\mathcal{N}$ of $\Delta^e(\mathcal{M}_1) \cup \Delta^e(\mathcal{M}_2)$. It follows in particular that $\mathcal{N}$ is a model of $T_1 \cup T_2 \cup T_0^*$ and that there are Σ_i-embeddings $h_i : \mathcal{M}_i \longrightarrow \mathcal{N}$. Notice that for $b \in \mathcal{M}$, we must have $h_1(b) = h_2(b)$;[14] to show that $\mathcal{N}$ is a $\Sigma_1 \cup \Sigma_2$-extension of $\mathcal{M}$ observe that for every n-ary Σ_i-function symbol f and for every n-tuple $\underline{b}$ of elements from the support of $\mathcal{M}$, we have $f^{\mathcal{M}}(\underline{b}) = f^{\mathcal{M}_i}(\underline{b}) = f^{\mathcal{N}}(h_i(\underline{b}))$; analogously, for every n-ary Σ_i-predicate symbol P, we have $\mathcal{M} \models P(\underline{b})$ iff $\mathcal{M}_i \models P(\underline{b})$ iff $\mathcal{N} \models P(h_i(\underline{b}))$. Thus, the interpretation of function and predicate symbols of $\Sigma_1 \cup \Sigma_2$ in $\mathcal{N}$ is just an extension of their interpretation in $\mathcal{M}$, which is another way of saying that $\mathcal{N}$ is a $\Sigma_1 \cup \Sigma_2$-extension of $\mathcal{M}$. $\dashv$

6 Applications

We say that a Σ_0-universal theory T_0 is *locally finite* iff Σ_0 is finite and for every finite set $\underline{a}$ of new free constants, there are finitely many $\Sigma_0^{\underline{a}}$-ground terms $t_1, \dots, t_{k_{\underline{a}}}$ such that for

[14]Recall that $\mathcal{N}$ is a model of $\Delta^e(\mathcal{M}_1) \cup \Delta^e(\mathcal{M}_2)$, hence it is a $\Sigma_1 \cup \Sigma_2 \cup \mathcal{M}_1 \cup \mathcal{M}_2$-structure; in this signature, there is a name $\bar{c}$ for every $c \in \mathcal{M}_i$ and h_i is defined by $h_i(c) = \bar{c}^{\mathcal{N}}$. For $b \in \mathcal{M}_1 \cap \mathcal{M}_2 = \mathcal{M}$, we have in particular that $\bar{b}^{\mathcal{N}}$ must be equal to $h_1(b)$ and $h_2(b)$ so that $h_1(b)$ and $h_2(b)$ are equal to each others.

every further $\Sigma_0^{\underline{a}}$-ground term u, we have $T_0 \models u = t_i$ (for some $i = 1, \ldots, k_{\underline{a}}$). As we are mainly dealing with computational aspects, we consider part of the definition the further request that such $t_1, \ldots, t_{k_{\underline{a}}}$ are effectively computable from $\underline{a}$. Examples of locally finite theories are the theory of graphs, of partial orders (more generally, any theory whose signature does not contain function symbols), of commutative idempotent monoids, of Boolean algebras, etc.

In a locally finite theory T_0, there are restricted *finite* classes which are representatives, up to T_0-equivalence, of the whole classes of $\Sigma_0^{\underline{a}}$-ground literals, clauses, quantifier free sentences, etc. (they are just the ground literals, clauses, quantifier free sentences, etc. containing only the above mentioned terms $t_1, \ldots, t_{k_{\underline{a}}}$). As it is evident that we can limit information exchange to ground positive clauses and quantifier-free sentences in that restricted class, both Theorems 4.1 and 4.2 yield *combined decision procedures for the clausal word problem* in $T_1 \cup T_2$ (in case T_0 is locally finite and in case the corresponding clausal word problems for T_1 and T_2 are separately decidable).

The procedure suggested by Theorem 4.1 is just a fair information exchange of positive ground $\Sigma_0^{\underline{a}}$-clauses, to be continued until the situation gets stable or until an inconsistency is detected. Notice that in case T_1, T_2 are Σ_0-convex theories, information exchange can be further limited to ground $\Sigma_0^{\underline{a}}$-atoms by Corollary 4.3.[15]

On the contrary, the procedure suggested by Theorem 4.2 (which is nothing but an interpolation theorem) identifies all ground $\Sigma_0^{\underline{a}}$-clauses which are logical consequences of $T_1 \cup \Gamma_1$ and check whether their conjunction is consistent with $T_2 \cup \Gamma_2$. There is a third possible (non deterministic) procedure, which is suggested in the introduction of the paper and which is justified directly by Lemma 5.2: as there are only finitely many Σ_0-structures generated by $\underline{a}$ which are models of T_0 (recall that such structures cannot have more than $k_{\underline{a}}$-elements), one simply guesses one of them and check whether its diagram is consistent with both $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$.

We give a first example to which the above outlined combined procedures apply.

Example. Let T_1 be rational linear arithmetic and let T_2 be the theory of total orders endowed with a strict monotonic function f.[16] We take as T_0 the theory of total orders (recall that its model completion T_0^* is the theory of dense total orders without endpoints). T_1 is known to be decidable [5]. We leave the reader the little exercise to prove that the clausal word problem for T_2 is decidable too: the relevant lemma to be proved

[15]This observation, as shown in [17] for the disjoint signatures case, may improve complexity bounds in certain significant situations.

[16]This means that f is subject to the axiom $\forall x \forall y\, (x < y \to f(x) < f(y))$.

consists in showing that any finite total order endowed with a *partial* strict monotonic function embeds into a model of T_2 (this is shown by successively inserting new points and by taking union in the limit). Once this lemma is proved, the satisfiability of a set Γ_2 of Σ_2^a-ground literals can be decided by a non deterministic guessing of such a finite total order endowed with a partial strict monotonic function. As $T_1 \supseteq T_0^*$, T_1 is certainly T_0-compatible. We only need to show that T_2 is T_0-compatible, by embedding each model $\mathcal{M}$ of T_2 into a model $\mathcal{M}'$ of $T_0^* \cup T_2$. It is sufficient to take as $\mathcal{M}'$ the lexicographic product of $\mathcal{M}$ with e.g. the poset of rational numbers[17] (the symbol f is interpreted by putting $f^{\mathcal{M}'}(b,m) = (f^{\mathcal{M}}(b), m)$, the embedding $\mathcal{M} \longrightarrow \mathcal{M}'$ is defined by associating with $b \in \mathcal{M}$ the pair $(b, 0)$). Thus our combination results apply and we obtain the decidability of the clausal word problem for rational linear arithmetic endowed with a strict monotonic function. It is not difficult to see that the complexity of this combined decision algorithm, if applied to satisfiability of the existential closure of a quantifier-free formula, lies in the NP-class (just adapt the arguments in [17]). $\dashv$

In order to explain our applications to fusion decidability in modal logic, we need to fix some terminology. We shall not directly introduce modal logic, rather we insist on its algebraic counterpart (this might look not very orthodox, but makes things more simple for the purposes of this paper). A *modal algebra* is just a Boolean algebra $\mathcal{B} = \langle B, \cap, 1, \cup, 0, (-)' \rangle$ endowed with an hemimorphism $\Box$ (a hemimorphism is a function preserving only binary meets and the top element). Hemomorphisms are also called modal (necessity) operators. The modal operator $\Box$ is said to be *transitive* iff the identity $\Box a \leq \Box\Box a$ holds for every $a \in B$.[18]

Let now Σ_1 be the signature of Boolean algebras augmented with a unary function symbol $\Box_1$ and let Σ_2 be the signature of Boolean algebras augmented with a unary function symbol $\Box_2$. T_1 is the equational theory of a variety V_1 of modal algebras and T_2 is the equational theory of another variety V_2 of modal algebras. For $i = 1, 2$, T_i is a universal Horn theory, hence it is Σ_i-convex: this means in particular that the solvability of the conditional word problem for T_i implies the solvability of the clausal word problem for T_i.

Proposition 6.1 *Let T_1, T_2 be as above; then the solvability of the conditional word problem for T_1 and T_2 implies the solvability of the conditional word problem for $T_1 \cup T_2$.*

[17] For density, observe that $(b,m) <_{lex} (b',m')$ implies $(b,m) <_{lex} (b,n) <_{lex} (b',m')$, where $n = \frac{m+m'}{2}$ in case $m < m'$ and $n = m+1$ otherwise (notice that in this last case, we must have $b < b'$).

[18] Recall that, in a Boolean algebra, $y \leq z$ is a shorthand for $y \cap z = y$.

Proof. As already mentioned, we have for free the solvability of the clausal word problem in T_1 and T_2; we take as T_0 the theory of Boolean algebras (which is locally finite and admits as a model completion the theory of atomless Boolean algebras): in order to apply our results, we simply need to show that T_1, T_2 are T_0-compatible. We do it for T_1. Let $\mathcal{M}$ be a model of T_1, we show how to embed it into a model $\mathcal{M}'$ of T_1 which is based on an atomless Boolean algebra: this is a well-known and rather trivial fact, which is used also in [24] as a side preliminary lemma. Instead of reporting the argument used in [24], we give a more direct one.[19] Define a sequence of models of T_1 by: $\mathcal{M}_0 := \mathcal{M}$, $\mathcal{M}_{k+1} := \mathcal{M}_k \times \mathcal{M}_k$; define also embeddings $\delta_k : \mathcal{M}_k \longrightarrow \mathcal{M}_{k+1}$ by $\delta_k(a) := \langle a, a \rangle$. Now take as $\mathcal{M}'$ the union (more precisely, the inductive limit) of this chain: clearly $\mathcal{M}'$ is atomless as a Boolean algebra (no non-zero element is minimal in it as any $a \in \mathcal{M}_k$ gets identified with $\langle a, a \rangle = \langle a, 0 \rangle \cup \langle 0, a \rangle$ in $\mathcal{M}_{k+1}$). ⊣

Corollary 6.2 *Let T_1, T_2 be as above; if the modal operators $\Box_1, \Box_2$ are transitive, then the solvability of the word problem for T_1 and T_2 implies the solvability of the word problem for $T_1 \cup T_2$.*

Proof. This is simply due to the fact that, because of the transitivity of the modal operators, the solvability of the word problem in T_i implies the solvability of the conditional word problem in T_i $(i = 1, 2)$: in fact the T_i-validity of the quasi-equation

$$t_1 \neq 1 \vee \cdots \vee t_n \neq 1 \vee t = 1$$

is equivalent to the T_i-validity of the single equation

$$\Box_i^+ t_1 \cap \cdots \cap \Box_i^+ t_n \leq t,$$

(where $\Box_i^+ t_k$ is $t_k \cap \Box_i t$) because of elementary facts concerning congruences in modal algebras. ⊣

We underline that the last corollary, once read in terms of logics, means exactly fusion decidability for normal extensions of $K4$. Although this does not entirely cover Wolter's fusion decidability results [24],[20] it puts some substantial part of them into the appropriate

[19] We thank L. Santocanale for suggesting this argument.

[20] The problem does not concern generalizations of Proposition 6.1 to logics containing many (also non-normal) modal operators (such generalizations are trivial and we omit them just for the sake of simplicity). The point is that in Corollary 6.2 we had to assume transitivity to get solvability of conditional word problems from solvability of word problems. Omitting transitivity, Corollary 6.2 still holds [24], however it might be problematic to get this further result in our context, because our combination method needs a preliminary purification step, leading by itself to a conditional word problem.

general combination context. For new results (based on a refinement of the combination schema explained in this section) concerning fusion of modal logics sharing a universal modality and nominals, see [10].

7 Combination in a saturation-based framework

The results in section 4 suffer of a main drawback: in order to take advantage from Theorems 4.1, 4.2, one should employ provers that are able to *enumerate* the $\Sigma_0^{\underline{a}}$-ground quantifier-free consequences of $T_i \cup \Gamma_i$, whereas it is well-known that existing provers (mainly based on a saturation process) do not do that. This is due to the difference between refutational completeness and completeness tout court: completeness is not considered as a desirable property for efficiency reasons. Of course the problem could be circumvented simply by testing for inconsistency the sets of sentences $T_i \cup \Gamma_i \cup \{\neg\phi\}$ (where ϕ is a quantifier-free ground $\Sigma_0^{\underline{a}}$-formula) in order to check that ϕ is a logical consequence of $T_i \cup \Gamma_i$ (this should be done for each ϕ incrementally, in order to get the desired enumeration). It goes without saying that this method is highly infeasible and innatural, so that a different solution must be found in order to take a real advantage from the ideas explained in section 4 (expecially whenever local finiteness of T_0 fails).

What we have in mind is a real integration of a prover dealing with $T \cup \Gamma_1$ and of a prover dealing with $T \cup \Gamma_2$: each of them pursues its own original saturation process, with a concurrent sharing of the $\Sigma_0^{\underline{a}}$-clauses which are run-time deduced. Of course, it is assumed that the two provers are based on homogeneous saturation methods (e.g. both on resolution, paramodulation, superposition calculi, etc.). This idea can be formalized by using the notion of a pure inference: an inference among $(\Sigma_1 \cup \Sigma_2)^{\underline{a}}$-clauses

$$\frac{C_1, \dots, C_n}{C}$$

is *pure* iff there is $i = 1, 2$ such that all the clauses $C_1, \dots, C_n, C$ are $\Sigma_i^{\underline{a}}$-clauses. Similarly, a deduction is pure iff all inferences in it are pure. Usually pure deductions are not able to detect inconsistency of (the skolemization of) sets of sentences like $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$, however we shall see that this may happen when the T_0-compatibility conditions are satisfied.

In order to realize this program, we first need to skolemize the theories T_1, T_2, thus passing to theories T_1^{sk}, T_2^{sk} in extended signatures $\Sigma_1^{sk}, \Sigma_2^{sk}$; Skolem functions will *not* be considered shared symbols, hence we still have that $\Sigma_0 = \Sigma_1^{sk} \cap \Sigma_2^{sk}$. The first problem we meet is the following: if T_i is T_0-compatible, is T_i^{sk} still T_0-compatible? We do not have a general answer for that, however there is a relevant case in which the answer is affirmative:[21]

Lemma 7.1 *Let T be a Σ-theory which is compatible with respect to a Σ_0-theory T_0 (here Σ_0 is a subsignature of Σ). If the axioms of T are all $\forall\exists$-sentences, then T^{sk} is T_0-compatible too.*

Proof. Let $\mathcal{M}$ be a model of T^{sk}; we want to embed it into a model of $T^{sk} \cup T_0^*$. As T is T_0-compatible, clearly the Σ-reduct of $\mathcal{M}$ is embeddable into a model $\mathcal{M}'$ of $T \cup T_0^*$. We only have extend the interpretation of the additional Skolem functions from $\mathcal{M}$ to $\mathcal{M}'$. Let $f_1, \dots, f_m$ be the n-ary Skolem functions coming from the T-axiom[22]

$$\forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \, \phi(x_1, \dots, x_n, y_1, \dots, y_m).$$

Pick $a_1, \dots, a_n$ from $\mathcal{M}'$; as $\mathcal{M}' \models T$, there are $b_1, \dots, b_m$ such that the quantifier-free ground formula $\phi(a_1, \dots, a_n, b_1, \dots, b_m)$ is true in $\mathcal{M}'$. In case $a_1, \dots, a_n$ all belong to the support of $\mathcal{M}$, we choose b_i to be precisely $f_i^{\mathcal{M}}(a_1, \dots, a_n)$ (this is possible because ϕ is quantifier-free, hence $\mathcal{M} \models \phi(a_1, \dots, a_n, b_1, \dots, b_m)$ entails $\mathcal{M}' \models \phi(a_1, \dots, a_n, b_1, \dots, b_m)$ in case the b_i's are from the support of $\mathcal{M}$). We put $f_i^{\mathcal{M}'}(a_1, \dots, a_n)$ equal to b_i and we are done. $\dashv$

The previous Lemma motivates the following extra assumption (in addition to those from section 4):

Assumption (IV). *T_1, T_2 are axiomatized by $\forall\exists$-sentences; T_1^{sk}, T_2^{sk} are their skolemizations.*

We take into consideration here the *superposition calculus* $\mathcal{I}$ (see [4], [15]). As commonly done within paramodulation-based theorem proving, we assume that identity is the only predicate[23] and we treat it symmetrically (this means that the equation $s = t$ is

[21] This case is important: recall from section 2 that T_0^* must have a set of $\forall\exists$-axioms, hence the cases in which T_0-compatibility is trivially ensured by the fact that the axioms for T_0^* are included in the axioms of T_i are not compromized by the skolemization process (provided the remaining axioms for T_i are also $\forall\exists$-sentences).

[22] Recall that this axiom is replaced in T^{sk} by the axiom

$$\forall x_1 \cdots \forall x_n \, \phi(x_1, \dots, x_n, f_1(x_1, \dots, x_n), \dots, f_m(x_1, \dots, x_n))$$

containing the new Skolem function $f_1, \dots, f_m$.

[23] Atomic formulas $P(t_1, \dots, t_n)$ are seen as equations $p(t_1, \dots, t_n) = \top$. For this and other similar details, see [15].

identified with the multiset $\{s, t\}$). Clauses[24] like

$$t_1 = u_1, \ldots, t_n = u_n \Rightarrow t'_1 = u'_1, \ldots, t'_m = u'_m$$

are often identified (e.g. for questions concerning orderings) with the multiset of multisets

$$\{\{t_1, t_1, u_1, u_1\}, \ldots, \{t_n, t_n, u_n, u_n\}, \{t'_1, u'_1\}, \ldots, \{t'_m, u'_m\}\}.$$

We fix a *lexicographic path ordering*[25] induced by a total precedence on the symbols of $\Sigma_1^{sk} \cup \Sigma_2^{sk} \cup \{\underline{a}\}$; assuming for simplicity that our signatures are finite, this induces a reduction ordering $>$ which is total on ground terms.[26] We give to symbols in $\Sigma_0^{\underline{a}}$ *lower precedence* than to symbols in $\Sigma_1^{sk} \setminus \Sigma_0$ and in $\Sigma_2^{sk} \setminus \Sigma_0$. This is essential: as a consequence, ground $\Sigma_0^{\underline{a}}$-clauses will be smaller in the twofold multiset extension of $>$ than all ground clauses containing a proper Σ_1 or Σ_2-symbol.

We recall below the inference rules of the superposition calculus $\mathcal{I}$: the application of each rule is subject to the satisfiability of certain *ordering constraints*, which are essential for the efficiency of the calculus; we do not report such constraints here for the sake of simplicity, we just remind that, roughly speaking, they *restrict the application of the rules to maximal terms/equations* in the respective clauses (see again [15] for details):

$$\frac{\Gamma' \Rightarrow \Delta', l = r \qquad \Gamma \Rightarrow s = t, \Delta}{\Gamma\mu, \Gamma'\mu \Rightarrow s\mu[r\mu]_p = t\mu, \Delta\mu, \Delta'\mu}$$

$$\frac{\Gamma' \Rightarrow \Delta', l = r \qquad \Gamma, s = t \Rightarrow \Delta}{\Gamma\mu, \Gamma'\mu, s\mu[r\mu]_p = t\mu \Rightarrow \Delta\mu, \Delta'\mu}$$

$$\frac{s = t, \Gamma \Rightarrow \Delta}{\Gamma\mu \Rightarrow \Delta\mu}$$

$$\frac{\Gamma \Rightarrow s = s', t = t', \Delta}{\Gamma\mu, s'\mu = t'\mu \Rightarrow s\mu = t'\mu, \Delta\mu}$$

[24] In this section, we use sequent notation for clauses.

[25] It is not clear whether the results explained in this section hold in case a Knuth-Bendix ordering is adopted.

[26] Our term-rewriting terminology is standard (see e.g. [1]).

where, in the first two rules, the premises do not share any variable and $s_{|p}$ is not a variable (the substitution μ is supposed to be a most general unifier of $\{s_{|p}, l\}$ in the first two rules and of $\{s, t\}$ in the last two rules).

Theorem 7.2 *In the above assumptions (I)-(IV), the set of sentences $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent iff there is a pure $\mathcal{I}$-derivation of the empty clause from $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$.*

Proof. We know that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is consistent iff so is $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$. Suppose that there is no pure $\mathcal{I}$-derivation of the empty clause from $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$. The saturation process with respect to pure $\mathcal{I}$-inferences produces a set of clauses Δ that can be represented as the union of two sets Δ_1, Δ_2 of $\Sigma_1^{sk} \cup \{\underline{a}\}$ and of $\Sigma_2^{sk} \cup \{\underline{a}\}$-clauses, respectively. Such sets can be taken so that they contain *the same* $\Sigma_0^{\underline{a}}$-clauses (otherwise said, $\Sigma_0^{\underline{a}}$-clauses from Δ are put in both of them); moreover, each of them is saturated with respect to $\mathcal{I}$-inferences from $T_i^{sk} \cup \Gamma_i$. Let $\Delta_1^{gr}, \Delta_2^{gr}$ be the sets of ground instances of clauses in Δ_1 and Δ_2. Following the model generation procedure explained in [15], we can define, by transfinite induction on ground clauses ordering, two convergent ground rewrite sytems R_1, R_2 whose associate normal Herbrand models $\mathcal{H}_1, \mathcal{H}_2$ are models of Δ_1, Δ_2 (hence also of $T_1^{sk} \cup \Gamma_1$ and of $T_2^{sk} \cup \Gamma_2$), respectively. The definition of R^i $(i = 1, 2)$ is as follows. Say that a $\Sigma_i^{sk} \cup \{\underline{a}\}$-ground clause C of the kind

$$\Gamma \Rightarrow l = r, \Delta$$

is *productive* and that that $Gen^i(C) = \{l \to r\}$ iff the following conditions are satisfied:

- $R_C^i \not\models C$;
- $l > r$, $l > u$ (for all u occurring in Γ), $\{l, r\} > \{u, v\}$ (for all equation $u = v$ occurring in Δ);
- l is not reducible by R_C^i;
- $R_C^i \not\models r = t'$, for all equation of the kind $l = t'$ occurring in Δ

(here R_C^i is the rewrite system $\bigcup_{D \in \Delta_i^{gr}, D < C} Gen^i(D)$). If C is not productive, put $Gen^i(C) = \emptyset$. Finally take R^i to be $\bigcup_{C \in \Delta_i^{gr}} Gen^i(C)$.

From the construction, it is clear that $\mathcal{H}_1, \mathcal{H}_2$ have the same Σ_0-substructure generated by $\underline{a}$ (up to a Σ_0-isomorphism): in fact, since ground $\Sigma_0^{\underline{a}}$-clauses/terms are smaller than clauses/terms containing a proper Σ_1 or Σ_2-symbol, it is immediate to prove by transfinite induction that a $\Sigma_0^{\underline{a}}$-term can only be reduced by a rule produced by a $\Sigma_0^{\underline{a}}$-clause and that, if C is such a clause, then $Gen^1(C) = Gen^2(C)$. Thus, R^1 and R^2 reduce a ground $\Sigma_0^{\underline{a}}$-term to the same normal form and this implies that $\mathcal{H}_1, \mathcal{H}_2$ share the same Σ_0-substructure $\mathcal{A}$ generated by $\underline{a}$.

By Lemma 7.1, T_1^{sk} and T_2^{sk} are both T_0-compatible, hence we can apply Lemma 5.2 to them and to the models $\mathcal{H}_1, \mathcal{H}_2$, thus getting a model $\mathcal{M}$ of $T_1^{sk} \cup T_2^{sk}$ and two $\Sigma_i^{sk} \cup \mathcal{A}$-embeddings $\mathcal{H}_i \longrightarrow \mathcal{M}$. As ground literals are preserved by embeddings, $\mathcal{M}$ is also a model of $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$, as desired. $\dashv$

Before concluding this section, we shall provide an example in which the assumptions of Theorem 7.2 are satisfied and an example in which such assumptions fail.

Example. Let T_1, T_2 be both the theory of Boolean algebras; we assume that symbols of the bounded distributivity lattice language (namely $\cap, \cup, 0, 1$) are shared but that the two complements n_1, n_2 are not. We want to prove that

$$T_1 \cup T_2 \models \forall x (n_1(x) = n_2(x)).$$

If we take T_0 to be the theory of bounded distributive lattices (i.e. of distributive lattices with 0 and 1), we see that T_1, T_2 are T_0-compatible.[27] Skolemization and purification give for instance the two sets of literals $\Gamma_1 = \{a = n_1(c), a \neq b\}$ and $\Gamma_2 = \{b = n_2(c), a \neq b\}$. A pure $\mathcal{I}$-refutation exists: the prover SPASS produces a pure $\mathcal{I}$-refutation consisting on 28 steps. However, the system is not programmed in order to avoid impure inferences, so that, during saturation, it impurely derives also (useless) 'mixed' clauses containing both n_1 and n_2. One of them, namely the atom $b \cap n_1(n_2(a)) = b$, is also selected as a given clause. $\dashv$

Example. Let T_1 be the theory of Boolean algebras and let T_2 be the theory of pseudocomplemented distributive lattices; these are bounded distributive lattices endowed with a unary operator $(-)^*$ satisfying the condition

$$\forall x \forall y \ (x \cap y = 0 \leftrightarrow y \leq x^*)$$

$(-)^*$ expresses the properties of intuitionistic negation, but in the union theory $T_1 \cup T_2$ such operator collapses into the classical complement. This is evident, because the axiom for $(-)^*$ implies that x^* is the supremum of the set of the elements y such that $x \cap y = 0$: as the Boolean complement of x enjoys the same property, clearly it must coincide with

[27] The model completion of T_0 is the theory of atomless Boolean algebras (formulated without the complement in the language): in fact, the theory of atomless Boolean algebras can eliminate quantifiers even without having the complement in the language (an atomic formula can be rewritten as the conjunction of equations of the kind $x'_1 \cup \cdots \cup x'_n \cup y_1 \cup \cdots \cup y_m = 1$ and these equations can be written without complement as $x_1 \cap \cdots \cap x_n \leq y_1 \cup \cdots \cup y_m$).

x^*. This means that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent, where Γ_1 is empty and Γ_2 is $\{(a^*)^* \neq a\}$. A SPASS refutation takes 43 lines and it is highly impure. In fact a pure refutation cannot exist: the Σ_0- (and even the Σ_0^a)-ground clauses deducible from either $T_1 \cup \Gamma_1$ or $T_2 \cup \Gamma_2$ are insufficient to detect inconsistency, because they are all subsumed by the three negative literals $0 \neq 1, a \neq 1, a \neq 0$. Notice that T_2 is not T_0-compatible. $\dashv$

8 Conclusions and related work

In this paper we have extended Nelson-Oppen combination procedure to the case of theories T_1, T_2 over non-disjoint signatures, in presence of compatibility conditions over a common universal subtheory T_0. The extension we proposed applies to examples of real interest giving, as shown in section 6, combined decidability in case T_0 is locally finite. Whenever T_0 is not locally finite, our method can be used in order to limit residue exchange (see section 4) or in order to forbid impure inferences in saturation-based theorem proving, thus yieldying restrictions on the search space during refutation derivations (see section 7).

It should be noticed that quantifier-elimination *plays only an indirect role* in the paper: in this sense, the existence of a model completion for a universal theory T_0 *guarantees a certain behaviour* in combination problems *by itself*, independently on how quantifier elimination in the model completion is established (this can be established also by semantic non constructive arguments, as largely exemplified in the model-theoretic literature). In principle, the quantifier elimination complexity/decidability *has nothing to do* with the complexity/decidability of our combination methods, simply because *quantifier elimination algorithms do not enter into them.* This is crucial, because most quantifier elimination algorithms are subject to heavy complexity lower bounds, which are often structural lower bounds for the decision of the elementary word problem in the corresponding theories (see the book [8]). In fact, there are instructive examples showing how complexity can grow from constraint satisfiability of (the existential closure of) conjunctions of literals to decision problems in full first-order language: the most striking one is the case of acyclic lists, where complexity grows from linear [18] to non-elementary [8], [18].

One may wonder how severe is the crucial condition of T_0-compatibility used in the paper: let us discuss it for a while. T_0-compatibility involves two aspects, namely the existence of a model completion T_0^* for T_0 and the embeddability of models of T_i into models of $T_i \cup T_0^*$. As we have shown in the examples, the existence of a model completion seems to be frequent for theories commonly used in software verification. On one side, numeric constraint theories often enjoy this property, in the sense that they eliminate quantifiers (thus being model completions of the theories axiomatized by their respective universal consequences). On the other side, acyclic binary lists might probably be the paradigm of situations arising in theories axiomatizing natural datatypes. Finally, notice that quantifier elimination strictly depends on the choice of the language: every theory trivially has quantifier elimination in an extended language with infinitely many definitional axioms, hence the problem of obtaining quantifier elimination seems to be mostly a problem of choosing a sufficiently rich bust still natural and manageable language.

The question concerning embeddability of models of T_i into models of $T_i \cup T_0^*$ looks more problematic, in the sense that it can fail in significant situations and, in addition, it does not look to be mechanizable. Further research is necessary on this point, however we underline that there is a relevant case in which the problem disappears. This is the case in which T_i is an extension of T_0^*: we have seen an example in section 6 where T_i is rational linear arithmetic and T_0 is the theory of linear orders. Another example is the theory of acyclic lists L (which coincides with L^*): any extension of the theory of acyclic lists with significant extra structures matches our requirements and the advantages of our method (limited residue exchange, elimination of impure inferences, etc) apply to all combinations of theories obtained in this way.

There have been many efforts in the literature trying to extend Nelson-Oppen combination method to theories sharing function and predicate symbols (different from equality). We leave apart interesting recent results [3], [9] on combined word problems because (as pointed out in [2]) they cannot be appropriately seen as generalizations of Nelson-Oppen combination procedure; rather, we concentrate here on [21] which is directly related to the subject.

The starting point of any attempt to generalize Nelson-Oppen procedure to the non-disjoint case should preliminarily answer the following question: what is the specific feature of the stable infiniteness requirement that we want to generalize? In the present paper we answered the question by saying that infinite models are just *existentially closed*[28] models of the pure theory of equality and based our further investigations on this observation. On the contrary, in [21] the authors notice that infinite models are just *free* models of the pure theory of equality with infinitely many generators. This leads to a completely different

[28]Let T_0 be a universal (for the sake of simplicity) theory. A model $\mathcal{N}$ is existentially closed for T_0 iff for every extension $\mathcal{M} \models T_0$ of $\mathcal{N}$, we have that every existential sentence (with parameters from $\mathcal{N}$) which holds in $\mathcal{M}$ holds in $\mathcal{N}$ too. It can be shown (see [6]) that T_0 has a model completion T_0^* iff the class of models of T_0 has the amalgamation property and the class $\mathbf{K}$ of models of T_0 which are existentially closed for T_0 is an elementary class (then, it turns out that K is precisely the class of models of T_0^*).

research direction, whose results can only very roughly be reported here.[29] Let T be a theory in the signature Σ and let Σ_0 be a finite subsignature of Σ; say that 'T is stably Σ_0-free (over a certain constraint language) iff every constraint (in the language) satisfiable in T is satisfiable in a model of T whose Σ_0-reduct is a free structure with infinitely many generators' (loc. cit. p.296). The main notion in [21] is that of N-O combinability of two theories T_1, T_2 over signatures Σ_1, Σ_2 (admitting Σ_0 as a common subsignature): N-O combinability guarantees, in essence, that 'the satisfiability in a theory $T_1 \cup T_2$ of the conjunction of two pure constraints $\phi_1 \wedge \phi_2$ can be reduced to the local satisfiability of ϕ_1 in T_1 and ϕ_2 in T_2 by adding to both formulas an appropriate Σ_0-restriction, a particular kind of first-order restriction on the free variables shared by ϕ_1 and ϕ_2' (loc. cit. p.295). One of the main results in [21] (Theorem 48, p.328) says that T_1, T_2 are N-O combinable provided they are both stably Σ_0-free and provided some additional conditions are satisfied.[30]

Apart from technical details, want we want to underline here is that the notion of infinitely generated free and of existentially closed structure are quite different and their coincidence for the pure theory of equality must be considered a rather exceptional fact. To see how the two notions can diverge consider the case of torsion-free abelian groups: now free abelian groups are torsion-free, but they are never divisible, whereas divisibility is just the necessary and sufficient condition for being existentially closed. Notice also that free structures may not exist for a theory, whereas every model of a universal theory T_0 embeds into a model which is existentially closed for T_0 (see [6]). These remarks should make evident that the generalization of the Nelson-Oppen procedure presented in [21] is quite different from that presented in this paper.

Before closing, we would like to remark that the idea (suggested in [19]) of using interpolation theorems in order to limit residue exchange in partial theory reasoning (whenever the background reasoner has to deal with combined theories) inspired some of the material presented in section 4 above. Notice however that we developed this intuition according to a philosophy which is in a sense opposite with respect to that of [19]: there the input theories T_1, T_2 were assumed to share all functions symbols (alien function symbols belonging to one theory being considered as free Skolem functions for the other), whereas we tried to keep function symbols separated too, as much as possible. This is essential in our context, because otherwise e.g. local finiteness of the common subtheory

[29] These results are formulated as constraint satisfiability results, because the main focus on [21] is on constraint satisfiability problems.

[30] These conditions are the collapse-free requirement for the theory of Σ_0-atoms entailed by T_1, T_2 and a stable Σ_0-freeness requirement for the union theory.

T_0 would be lost (and decidability of the combined problems presented in section 6 would not be achieved as a consequence).

9 Appendix A: quantifier elimination for acyclic lists

In this Appendix we prove the result concerning quantifier elimination for the theory of acyclic lists which is mentioned in section 2.

We recall that in order to show that a given theory T admits quantifier elimination is sufficient to prove that for every primitive formula $\exists x_0 \cdots \exists x_n \phi$ there is a quantifier-free formula ϕ' such that

$$T \models \exists x_0 \cdots \exists x_n \phi \leftrightarrow \phi'$$

(a formula $\exists x_0 \cdots \exists x_n \phi$ is said to be *primitive* iff ϕ is a conjunction of literals).[31] Once quantifiers are eliminated from primitive formulas, it is easy to eliminate them from arbitrary formulas, by using prenex normal forms (and the interdefinability of existential and universal quantifiers). The matrix ϕ of our primitive formula may contain additional free variables other than $x_0, \dots, x_n$; without loss, we replace them by distinct new free constants. Consequently, we consider below a language enriched by free constants and try to prove that all primitive *sentences* in that language are equivalent to quantifier free *sentences* in that language (this gives the quantifier elimination results, once free constants are replaced back by variables). Like in section 7, we shall treat identity symmetrically, meaning that $t_1 = t_2$ (resp. $t_1 \neq t_2$) may stand also for $t_2 = t_1$ (resp. $t_2 \neq t_1$). We shall use the notation $t(u/v)$ in order to indicate the term obtained from t by replacing in it all occurrences of the subterm v by the term u.

We recall from section 2 that the signature of the theory L consists on two unary function symbols car, cdr and one binary function symbol $cons$. We say that a term t is a component term (briefly, *κ-term*) iff it is built up from variables and constants by using finitely many applications of the unary function symbols car, cdr. A pairing term (briefly, a *π-term*) is a term built up from variables and constants by using the binary function symbol $cons$ only. A κ-term or a π-term is *proper* iff it is not a variable or a constant. An occurrence p of a term t in an expression E is a *π-occurrence* iff no symbol car, cdr occurs below it (namely, iff for every $q \leq p$, the root symbol of $E_{|q}$ is not car, cdr).

We recall the axioms of L:

[31] Usually, it is assumed $n = 0$ in the definition of a primitive formula, however our quantifier elimination procedure is not direct and it can introduce, during execution, new bounded variables. This is why we give a modified definition of a primitive formula, allowing more than one existential quantifier in the prefix.

$$
\begin{aligned}
&cons(car(x), cdr(x)) = x \\
&car(cons(x, y)) = x \\
&cdr(cons(x, y)) = y \\
&x \neq t(x)
\end{aligned}
$$

where t is a proper κ-term involving the variable x.

Lemma 9.1 *Let u be a term containing a π-occurrence of the variable x and let t be any term. Then there is a κ-term $v(y)$ such that*

$$L \models t = u \to x = v(t/y).$$

(the κ-term v is proper if $u \not\equiv x$).

Proof. This is evident, as e.g. $t = cons(u_1, cons(x, u_2))$ is equivalent to $car(t) = u_1 \wedge cdr(car(t)) = x \wedge cdr(cdr(t)) = u_2$, so it implies $x = cdr(car(t))$. A formal argument works by induction on the length of the π-occurrence p of x in u. If $p = \epsilon$, take v to be the variable y; otherwise (as u must be of the kind $cons(u_1, u_2)$), we have $p = 1p'$ or $p = 2p'$, so that $t = u$ implies both $car(t) = u_1$ and $cdr(t) = u_2$ and induction hypothesis applies to one of the latter equations. $\dashv$

We first need a very special case of quantifier elimination:

Lemma 9.2 *Let x_0 be a variable and let $y_1, \ldots, y_m$ be further variables (distinct from x_0 but not necessarily distinct from each other); also let $t_1, \ldots, t_k$ be terms not containing any occurrence of x_0 and let $u_1, \ldots, u_m$ be terms containing at least one π-occurrence of x_0. Then the formula*

$$(*) \qquad \exists x_0 (x_0 \neq t_1 \wedge \cdots \wedge x_0 \neq t_k \wedge y_1 \neq u_1 \wedge \cdots \wedge y_m \neq u_m)$$

is a logical consequence of L.

Proof. The formula $(*)$ is logically implied (under the axioms of L) by a formula of the kind

$$(+) \qquad \exists x_0 \, (x_0 \neq t_1 \wedge \cdots \wedge x_0 \neq t_k \wedge x_0 \neq v_1(y_1) \wedge \cdots \wedge x_0 \neq v_m(y_m))$$

where $v_1, \ldots, v_m$ are κ-terms (see Lemma 9.1). As x_0 does not appear neither in the t_j's nor in the v_i's, it is clear that $(+)$ is a logical consequence of L (because all models of L are infinite, as explained in section 2). $\dashv$

We are now ready for the whole quantifier elimination procedure:

Theorem 9.3 *L enjoys quantifier elimination.*

Proof. We eliminate quantifiers from a primitive sentence

$$\exists x_0 \cdots \exists x_n \phi$$

in four steps.

Step 1. The first three axioms of L give a convergent rewrite system: we normalize all terms in ϕ with respect to it. As an effect of this normalization, notice that if any term $\phi_{|p}$ has root symbol *car* or *cdr*, then it is a proper κ-term.

Step 2. If there is in ϕ a subterm of the kind $car(x_i)$ or $cdr(x_i)$, $\exists x_0 \cdots \exists x_n \phi$ is modified as follows:

$$\exists y_1 \exists y_2 \exists x_0 \cdots \exists x_n (x_i = cons(y_1, y_2) \wedge \phi(y_1/car(x_i), y_2/cdr(x_i)).$$

This transformation is repeated until all proper κ-terms are ground (this loop terminates because each single transformation decreases the number of the symbols car, cdr within ϕ). Notice also that, after this step, if any term $\phi_{|p}$ has root symbol *car* or *cdr*, then it is a proper and *ground* κ-term (we leave the reader to check that this fact is permanent during the next step transformations too). Thus, *all occurrences of variables are π-occurrences.*

Step 3. This is the main step. Apply non-deterministically the following transformations until possible:

3.1 Replace $cons(t, u) = cons(t', u')$ by $t = t' \wedge u = u'$.

3.2 Replace $cons(t, u) \neq cons(t', u')$ by $t \neq t' \vee u \neq u'$ (then distribute the disjunction over the conjunctions and the existential quantifiers in front of the formula and go on separately with each disjunct).

3.3 If v is a ground κ-term, replace $cons(t, u) = v$ by $t = car(v) \wedge u = cdr(v)$.

3.4 If v is a ground κ-term, replace $cons(t, u) \neq v$ by $t \neq car(v) \vee u \neq cdr(v)$ (then distribute the disjunction over the conjunctions and the existential quantifiers in front of the formula and go on separately with each disjunct).

3.5 Remove equations like $t = t$.

3.6 Stop, with output $\perp$, in case there is an equation like $t \neq t$.

3.7 If there is an equation like $x_i = t$, where x_i does not occur in t, remove it and replace x_i by t everywhere else (the quantifier $\exists x_i$ in front of the formula is deleted).

3.8 If there is an equation like $x_i = t$ (with $t \not\equiv x_i$), where x_i occurs in t, stop with output $\perp$.[32]

3.9 If there is a disequation like $x_i \neq t$ (with $t \not\equiv x_i$), where x_i occurs in t, remove it.[33]

The termination of this step can be shown by using a lexicographic product of triples of natural numbers $\langle N_1, N_2, N_3 \rangle$, where N_1 is the number of the quantified variables, N_2 is the number of occurrences of the symbol *cons* and N_3 is the number of literals.

Step 4. Keep only ground literals (this is the output quantifier-free sentence). This step is justified by Lemma 9.2. In fact, if $\exists x_0 \cdots \exists x_n \phi$ is our primitive sentence after Steps 1,2,3, then each literal in it is either ground or of the kind $x_i \neq t$ (for a term t in which x_i does not occur). Consequently, we can rearrange our primitive sentence in the form

$$\exists x_1 \cdots \exists x_n (\phi_1 \wedge \exists x_0 \phi_2)$$

where ϕ_1 does not contain x_0 and $\exists x_0 \phi_2$ matches the conditions in the Lemma 9.2: this shows that $\exists x_0 \phi_2$ can simply be removed. Doing that repeatedly, we realize that $\exists x_0 \cdots \exists x_n \phi$ is actually equivalent to the conjunction of its ground literals. $\dashv$

The above procedure in highly unfeasible if it is adopted as a decision algorithm for the elementary word problem: in fact, there are various points in it (e.g. 3.2 and 3.7, just to mention a couple of them) which can produce an exponential space explosion in case they are exhaustively applied. This is expecially harmful, because the whole procedure should be repeated as many times as there are alternating blocks of quantifiers in the prenex normal form of the input formula from which we want to eliminate quantifiers (notice also that the analysis of a universal block requires a preliminary rewriting of $\forall \cdots \forall$ into $\neg \exists \cdots \exists \neg$, followed by a DNF-conversion). This makes clear that the above algorithm does not have an elementary recursive complexity bound (as expected, from the corresponding absolute lower bound in [8]). However, see the remarks in section 8 for the right interpretation of the role of quantifier elimination in our combination context.

[32]As all occurrences of x_i are π-occurrences and as there is at least one such π-occurrence of x_i in t, the equation $x_i = t$ implies (by Lemma 9.1) an inconsistent equation of the kind $x_i = v(x_i)$, where v is a proper κ-term.

[33]Such a disequation is a logical consequence of L again by Lemma 9.1, given that all occurrences of x_i are π-occurrences.

10 Appendix B: submodel-completenss

The definition of a model completion of a theory T we used in the paper is specific for the case in which T is universal; here we recall some basic material, thus showing the reduction to the standard definition (taken e.g. from [22]).

We say that a Σ-theory T is *submodel-complete* iff given two models $\mathcal{M}_1, \mathcal{M}_2$ of T and given a common Σ-substructure $\mathcal{A}$ of them, we have that $\mathcal{M}_1$ is elementarily equivalent to $\mathcal{M}_2$ as a $\Sigma^{\mathcal{A}}$-structure. Taking in mind the diagram theorem, we may equivalently say that the theory T is submodel-complete iff $T \cup \Delta(\mathcal{A})$ is complete whenever $\mathcal{A}$ is a Σ-substructure of a model of T. We have that:

Proposition 10.1 *T admits elimination of quantifiers iff it is submodel-complete.*

Proof. One side is easy (see in any case the proof of Lemma 5.2). Suppose, for the other side, that T is submodel-complete and let $\phi(\underline{x})$ be an arbitrary formula. For new constants $\underline{a}$ consider the set of sentences

$$\Theta \ := \ T \cup \{\phi(\underline{a})\} \cup \{\neg\psi(\underline{a}) \,|\, \psi \text{ is quantifier free and } T \models \psi(\underline{a}) \to \phi(\underline{a})\}.$$

If Θ is inconsistent, then we have $T \models \phi(\underline{a}) \to \psi_1(\underline{a}) \vee \cdots \vee \psi_n(\underline{a})$ for quantifier-free ψ_i implying ϕ, so that we have

$$T \models \phi(\underline{x}) \leftrightarrow \psi_1(\underline{x}) \vee \cdots \vee \psi_n(\underline{x})$$

yielding the fact that T admits quantifier elimination. Consequently it suffices to show that Θ cannot be consistent. Suppose it is and let $\mathcal{M}$ be a model of it. Let $\mathcal{A}$ be the substructure of $\mathcal{M}$ generated by the $\underline{a}$; we must have

$$T \cup \Delta(\mathcal{A}) \models \phi(\underline{a}),$$

because otherwise we would be able to build a model of T containing $\mathcal{A}$ as a substructure and falsifying $\phi(\underline{a})$ (which cannot be because T is submodel-complete and $\phi(\underline{a})$ is true in $\mathcal{M}$ being an element of Θ). This means that for some quantifier-free sentence $\psi(\underline{a})$ true in $\mathcal{A}$ we have that $T \models \psi(\underline{a}) \to \phi(\underline{a})$. According to the definition of Θ, $\neg\psi(\underline{a})$ is true in $\mathcal{M}$ and also in $\mathcal{A}$ (because it is quantifier-free), contradiction. $\dashv$

Proposition 10.2 *Let T be a universal Σ-theory and let T^* be a further Σ-theory extending T. We have that T^* is a model completion of T iff the following two conditions are satisfied: (i) every model of T embeds into a model of T^*; (ii) for every Σ-structure $\mathcal{A}$ which is a model of T, we have that $T^* \cup \Delta(\mathcal{A})$ is a complete $\Sigma^{\mathcal{A}}$-theory.*

Proof. Assume that T is universal, that $T^* \supseteq T$ and that (i) holds. What we need to show is that (ii) holds iff T^* has quantifier elimination. By Proposition 10.1, this amounts to show that a Σ-structure $\mathcal{A}$ is a model of T iff it is a substructure of a model of T^*. This is evident: on one side just use (i) and on the other side use the fact that T is a universal subtheory of T^*. $\dashv$

The reason why we did not use in the paper the equivalent (and more general) definition suggested by the above Proposition is that we wanted to emphasize the role of quantifier elimination, which has a more direct meaning within automated deduction tradition.

References

[1] Baader F, Nipkow T., *Term Rewriting and All That*, Cambridge University Press, (1998).

[2] Baader F., Tinelli C., *Deciding the Word Problem in the Union of Equational Theories*, Technical Report UIUCDCS-R-98-2073, Department of Computer Science, University of Illinois at Urbana-Champaign (1998).

[3] Baader F., Tinelli C., *Deciding the Word Problem in the Union of Equational Theories*, Information and Computation, 178, 2, pp.346-390 (2002).

[4] Bachmair L., Ganzinger H. *Equational Reasoning in Saturation-Based Theorem Proving*, in Bibel L., Schmitt P.H. (eds.) "Automated Deduction - A Basis for Applications", vol. I, pp. 353-397, Kluwer (1998).

[5] Bockmayr A., Weispfenning V., *Solving Numerical Constraints*, in Robinson A., Voronkov A., (eds.) "Handbook of Automated Reasoning", vol. I, Elsevier/MIT, pp. 751-842 (2001).

[6] Chang C.C., Keisler H.J., *Model Theory*, IIIrd edition, North Holland (1990).

[7] Cooper D. C., *Theorem Proving in Arithmetic without Multiplication*, in Meltzer B., Milchie D., (eds.) "Machine Intelligence 7", pp. 91-99 (1972).

[8] Ferrante J., Rackoff C.W., *The Computational Complexity of Logical Theories*, Springer Lecture Notes in Mathematics 718, (1979).

[9] Fiorentini C., Ghilardi S., *Combining Word Problems through Rewriting in Categories with Products*, Theoretical Computer Science, 294, pp.103-149 (2003).

[10] Ghilardi S., Santocanale L., *Algebraic and Model Theoretic Techniques for Fusion Decidability in Modal Logic*, in preparation, (2003).

[11] Ghilardi S., Zawadowski M., *Sheaves, Games and Model Completions*, Trends in Logic Series, Kluwer (2002).

[12] Janiciċ P., Bundy A., *A General Setting for Flexibly Combining and Augmenting Decision Procedures*, Journal of Automated Reasoning, 28, pp.257-305 (2002).

[13] MacIntyre A., *Model Completeness*, in Barwise J. (ed.), "Handbook of Mathematical Logic", North Holland, pp. 139-180 (1977).

[14] Nelson G., Oppen D., *Simplification by Cooperating Decision Procedures*, ACM Transactions on Programming Languages and Systems, 1(2), pp. 245-257 (1979).

[15] Nieuwhenhuis R., Rubio A., *Paramodulation-Based Theorem Proving*, in Robinson A., Voronkov A., (eds.) "Handbook of Automated Reasoning", vol. I, Elsevier/MIT, pp. 371-533 (2001).

[16] Oppen D., *A* $2^{2^{2^{pn}}}$*-Upper Bound on the Complexity of Presburger Arithmetic*, Journal of Computer and Systems Sciences, 16(3), pp. 323-332 (1978).

[17] Oppen D., *Complexity, Convexity and Combination of Theories*, Theoretical Computer Science, 12, pp. 291-302 (1980).

[18] Oppen D., *Reasoning about Recursively Defined Data Structures*, Journal of the ACM, 27, 3, pp. 403-411 (1980).

[19] Tinelli C., *Cooperation of Background Reasoners in Theory Reasoning by Residue Sharing*, Journal of Automated Reasoning, 2003 (to appear).

[20] Tinelli C., Harandi M., *A New Correctness Proof of the Nelson-Oppen Combination Procedure*, in Baader F., Schulz K. (eds.) "1st International Workshop on Frontiers of Combining Systems (FroCos'96)", Applied Logic Series, vol. 3, Kluwer Academic Publishers, pp. 103-120 (1996).

[21] Tinelli C., Ringeissen C., *Unions of Non-Disjoint Theories and Combination of Satisfiability Procedures*, Theoretical Computer Science 290(1), pp.291-353, (2003).

[22] Wheeler W. H., *Model-Companions and Definability in Existentially Complete Structures*, Israel Journal of Mathematics, 25, pp.305-330 (1976).

[23] Weispfenning V., *The Complexity of Linear Problems in Fields*, Journal of Symbolic Computation, 5 (1-2), pp. 3-27 (1988).

[24] Wolter F., *Fusions of Modal Logics Revisited*, in Kracht M., De Rijke M., Wansing H., Zakharyaschev M. (eds.) "Advances in Modal Logic", CSLI, Stanford (1998).